

Hitachi NAS Platform

NAS Platform Integration with ELK

This document details the best practices for integrating HNAS with ELK (Elasticsearch).

MK-92HNAS086-01

May 2023

© 2021 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara Corporation (collectively, "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation at https://support.HitachiVantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara Corporation.

By using this software, you agree that you are responsible for:

- 1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and
- 2) Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

EXPORT CONTROLS - Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, and Licensee shall not export, or allow the export or re-export of, the Software, API, or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, Flash Copy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, BCPI™ and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or web site are properties of their respective owners.

Table of Contents

Preface	4
About this document.....	4
Document conventions.....	4
Intended audience.....	4
Accessing product downloads.....	4
Comments	4
Getting Help.....	5
Chapter 1: Overview.....	6
Chapter 2: Log collection	6
Log collection using Elastic Docker images	6
Modifying the hnas-syslog.conf file	6
Running Logstash in a container	7
Statistics Collections	8
HNASbeat YAML Configuration File	8
Statistic collection suggestions	9
Performance Envelope:	9
Basic Server Stats:	9
Cluster Throughput:	9
Cluster Redirections:	9
Disk Operations:	10
Protocol Operations:	10
FS Metadata Caches:	10
FS Sector Cache:	10
FSA Utilization:	10
Getting a List of Supported Statistics	10
NAS server setup	11
Syslog alert setup.....	11
Management auditing setup.....	11
Filesystem auditing setup.....	12
Testing the setup.....	12
Limitations/considerations.....	13

Preface

About this document

This document details the best practices for integrating HNAS with ELK (Elasticsearch). It assumes that you are familiar with the HNAS platform.

Document conventions

This document uses the following typographic convention:

Convention	Description
Bold	<ul style="list-style-type: none">Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK.Indicates emphasized words in list items.
<i>Italic</i>	Indicates a document title or emphasized words in text.
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>

Intended audience

This document is intended for customers and Hitachi Vantara partners who license and use NAS Platform.

Accessing product downloads

Product software, drivers, and firmware downloads are available on Hitachi Vantara Support Connect: <https://support.hitachivantara.com/>.

Log in and select Product Downloads to access the most current downloads, including updates that may have been made after the release of the product.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara.

Getting Help

[Hitachi Vantara Support Connect](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Chapter 1: Overview

The NAS server platform (Hitachi NAS platform and NAS module) can be integrated with Elasticsearch (<https://www.elastic.co/>). Alert and audit logs can be collected, and then analyzed using Kibana, which helps to visualize data. Elasticsearch is commonly referred to as the ELK stack or Elastic stack, which refers to Elasticsearch and associated components, which lets you reliably and securely take data from any source, in any format, and search, analyze, and visualize it in real time.

All instructions in this document assume that Elasticsearch and Kibana are already installed and working. They also assume that Docker is installed and available.

Chapter 2: Log collection

Log collection is achieved using a custom configuration file for Logstash, which is another component of the ELK stack, and provides a server-side data processing pipeline that ingests data from a multitude of sources simultaneously. The recommended approach is to use a containerized version of Logstash, downloaded from the Elastic website.

Chapter 3: Log collection using Elastic Docker images

Elastic provide Docker images to simplify deployment of their software. It's possible to use the Elastic provided Logstash Docker image with the Hitachi NAS specific config file, eliminating much of the complexity of installing Logstash.

All Elastic Docker images are available from <https://www.docker.elastic.co/>. Find the Logstash section and follow the instructions to download the image that corresponds to the version you require – either choose the latest release, or choose the one that matches the version of Elasticsearch already running locally.

Using version 6.6.0 as an example, run the following command (choose the oss version as it is smaller and none of the additional software included in the other image is required):

```
docker pull docker.elastic.co/logstash/logstash-oss:6.6.0
```

Once the image has been downloaded, set up the `hnas-syslog.conf` file before attempting to run the image.

Modifying the `hnas-syslog.conf` file

The only modification that should be made to the `hnas-syslog.conf` file is in the `elasticsearch` part of the `output` section. Supply the Elastic Search server host address and any credentials here. Do not make changes to any other sections of this file.

```
output {
  elasticsearch {
    hosts => ["http://<Elastic Search host IP address>:9200"]
  }
  stdout {
    codec => rubydebug
  }
}
```

```
}  
}
```

Refer to the following documentation for a full list of valid output parameters, and their syntax:

<https://www.elastic.co/guide/en/logstash/current/plugins-outputs-elasticsearch.html>

The `stdout` section causes all events received to be printed to the console (inside the container) and this is useful when testing that syslog events are being correctly received – this section can be removed once it has been established that log collection is working.

Running Logstash in a container

As syslog messages are sent on port 514, which is classed as a privilege port on Linux systems, and is generally only available to the root user, the `docker` command must be run using `sudo`, as follows. The configuration file tells syslog to listen on port 5140 instead of 514 because of this issue, so a mapping is required for both TCP and UDP when running the **docker** command. As the config file is also stored externally to the image, it must be mapped to within the running image with the `-v` option.

The `<folder>` parameter must specify a full path that contains the `hnas-syslog.conf` file - no other files should exist in the same folder.

```
sudo docker run -d --rm -p 514:5140/tcp -p 514:5140/udp \  
  -v <folder>:/usr/share/logstash/pipeline/ \  
  docker.elastic.co/logstash/logstash-oss:6.6.0
```

The **docker** command has the following parameters:

- `-d` detach and run in the background
- `--rm` automatically clean up the container and remove the file system when the container exits
- `-p` network port mapping between internal container port and external port on host machine
- `-v` mount folder to a specific location within the container

Note: For correct operation, the Logstash container image may not be run in an environment that makes use of Network Address Translation (NAT), such as Kubernetes. The source address of the NAS system sending events needs to reach Logstash to be recorded, and this does not happen if NAT is being used.

Chapter 4: Statistics Collections

HNAS statistics can be collected by a custom Hitachi NAS beat – in Elasticsearch term, a beat is lightweight single purpose data shipper. The beat can gather statistics from multiple NAS servers. The Hitachi NAS beat is supplied as a container image file and should be run in the background as follows. The only external information that needs to be supplied is the configuration file, which specifies which NAS servers to collect statistics from, and what statistics they are.

The Docker image must first be imported/loaded onto the machine that will be used to run the container:

```
docker load < hnasbeat-v1.0.tgz
```

Once the image has been imported, the following command can be used to run the new image in a container. The `<folder>` parameter must specify a full path that contains the `hnasbeat.yaml` configuration file (no other files should be present in the same folder) – see the next section on how to configure the `hnasbeat.yaml` configuration file.

```
docker run -d --rm -v <folder>:/config/ hitachi/hnasbeat:1.0
```

The above command starts the Hitachi NAS beat and detaches from the console, so it will run in the background. To help with troubleshooting, it can also be run in the foreground to provide additional debug output using the following command:

```
docker run --rm -v <folder>:/config/ hitachi/hnasbeat:1.0 -d ""
```

HNASbeat YAML Configuration File

The YAML configuration file for the HNAS beat needs to be of the format shown below. The `period` configures how often the statistics are gathered. 30 seconds is recommended as the default value, but if many statistics are to be collected, or there are many different servers, then increasing this value may be necessary to ensure that statistics are successfully collected.

You can configure multiple NAS systems. Specify each one under the `hosts` section, and the `address` value should be the IPv4 address for each physical node that needs to be monitored, as statistics are generally specific to a physical node. For monitoring a cluster, all nodes should be configured in the YAML file. The `stats` section contains a YAML formatted list of statistics that are to be collected from the specific address, and can be different for each system if required, but must be specified for each system – by default, no statistics are collected unless specified.

The following example collects statistics from two systems, every 30 seconds, and collects different statistics from each system.

```
hnasbeat:
  period: 30s
  hosts:
    - address: <NAS host IP address>
      stats:
        - Total Operations per Second
        - MMB Load (%)
```

```

- MFB Load (%)
- Ethernet Throughput RX (Mbps)
- Ethernet Throughput TX (Mbps)
- FibreChannel Throughput RX (Mbps)
- FibreChannel Throughput TX (Mbps)
- NVRAM waited allocs (#)
- address: <Second NAS host IP address>
  stats:
    - Total Operations per Second
    - MMB Load (%)
    - MFB Load (%)
output.elasticsearch:
  hosts: ["https://<Elastic Search host IP address>:9200"]

```

Statistic collection suggestions

The sections below cover some predefined suggestions of useful statistics to collect. It should be possible to copy and paste the text below directly into a YAML configuration file.

Performance Envelope:

```

- Total Operations per Second
- MMB Load (%)
- MFB Load (%)
- NVRAM waited allocs (#)
- Running Bossock Fibers (#)
- Running pi-tcp-sockets receive Fibers (#)

```

Basic Server Stats:

```

- Total Operations per Second
- MMB Load (%)
- MFB Load (%)
- Ethernet Throughput RX (Mbps)
- Ethernet Throughput TX (Mbps)
- FibreChannel Throughput RX (Mbps)
- FibreChannel Throughput TX (Mbps)
- NVRAM waited allocs (#)

```

Cluster Throughput:

```

- HSSI-C1 Throughput RX (Mbps)
- HSSI-C1 Throughput TX (Mbps)
- HSSI-C2 Throughput RX (Mbps)
- HSSI-C2 Throughput TX (Mbps)

```

Cluster Redirections:

```

- Total Operations Received per Second
- Total Forwarded Operations Received per Second
- Total Operations Forwarded per Second

```

Disk Operations:

- Current Virtual Disk Requests (#)
- Disk Read Latency (ms)
- Disk Write Latency (ms)
- Disk Stripe Write Latency (ms)

Protocol Operations:

- Protocol : iSCSI Operations per Second
- Protocol : NFS Operations per Second
- Protocol : SMB Operations per Second
- Protocol : SMB2 Operations per Second

FS Metadata Caches:

- wfile cache hits (%)
- wdir cache hits (%)
- wtree cache hits (%)
- obj_store root onode cache hits (%)
- obj_store non-ind-obj leaf onode cache hits (%)
- obj_store ind-obj leaf onode cache hits (%)

FS Sector Cache:

- Sector Cache Read Hits (%)
- Sector Cache ReadAhead Hits (%)
- Sector Cache Write Hits (%)

FSA Utilization:

- FSA Cache usage (%)
- Heap Usage (%)
- Running Bossock Fibers

Getting a List of Supported Statistics

As the statistics supported by the NAS sever are version specific, and depend on various configuration options, it's not possible to supply a full list here, but it is possible to ask the NAS server what statistics it supports (there are over 5000 different statistics available) as follows:

1. Telnet to the NAS server admin address
2. Connect to TCP port 11106

3. Type LIST.

The server lists all the currently available statistics. End the connection with QUIT. See the truncated example below:

```
root@lab:~# telnet <NAS server host IP address> 11106
Trying <NAS server host IP address>...
Connected to <NAS server host IP address>.
Escape character is '^]'.
LIST
0      Number of ssfs operations executed directly by the BOS Thread
1      sockets: New connections while under msgb stress
2      sockets: New connections denied while under msgb stress
3      sockets: New connections while under heap stress
4      sockets: New connections denied while under heap stress
....
6612   Reverse migration: wait for pool slot timed out (#)
6613   Reverse migration: wait for pool slot didn't time out (#)
.
QUIT
Connection closed by foreign host.
```

The statistics supported by the server are returned one per line and are prefixed by a number – use the text after the number and copy to the end of the line before inserting it into the YAML config file.

Chapter 5: NAS server setup

The following settings now need to be applied to the NAS server so that it can send its alert messages to the Logstash server. The following instructions use Command Line Interface (CLI) commands, as not all the functionality can be configured through the NAS Manager GUI.

Syslog alert setup

Configure the syslog event frequency and destination. The example below configures the NAS server to send all alerts, at all levels to the Logstash server (this is the recommended setting). Although fewer events can be configured to be sent out, it may prove less useful during analysis later.

```
alert-syslog set -s i -w i -i i
alert-syslog add <LogstashServerAddress>
```

There is also a variable to make sure that the syslog alerts come from the Admin IP address – this is more useful for single node setups, but can help with clustered setups too:

```
set syslog_and_snmptrap_come_from_admin_vnode true
```

Management auditing setup

The management audit logs messages need to be configured separately, using the following command:

```
audit-mgmt-log-server-add <LogstashServerAddress>
```

Filesystem auditing setup

Filesystem audit log messages also need to be configured separately and are normally accessed using other specialist tools. But it is possible to send them through syslog and ingest them into Elasticsearch.

Only a single syslog destination can be configured for this functionality.

```
audit-syslog add <LogstashServerAddress>
```

Note: Filesystem audit messages are only generated when the filesystem objects are configured accordingly – that configuration is outside the scope of this document.

Testing the setup

Once the Logstash server has been configured and the NAS server set up to send event and audit logs to the Logstash server, test the setup to make sure all types of event are sent to Elasticsearch.

Generate some audit events:

- Login to the NAS server either through the NAS Manager or SSC – this action generates various authentication events.

Generate a test event log entry using the following CLI command:

- `alert-send-test`

Below is an example of an audited login through SSC, and the Syslog test event - both are shown in Kibana. The `address` field can be used to check which device events are from:

```
▶ 2019-02-25 14:20:20.600 type: syslog severity: Information event: Test event. address: 172.16.19.240 host: 172.16.19.240
@timestamp: 2019-02-25 14:20:20.600 syslog_severity: informational syslog_severity_code: 6 @version: 1
eventID: 9999 syslog_facility: user-level syslog_facility_code: 1 _id: 6w8H3WkBVRIye0jXKySB _type: doc
_index: logstash-2019.02.25 _score: -

▶ 2019-02-25 14:20:16.778 type: syslog startTime: 1551104416774 severity: 2 deviceHostName: accept-m2.uktest.dev.bluearc.com host: acce
pt-m2-evs1.uktest.dev.bluearc.com @timestamp: 2019-02-25 14:20:16.778 syslog_severity: notice
destinationUserPrivileges: SUPERVISOR eventOutcome: success syslog_severity_code: 5 destinationUserName: super
visor deviceVersion: 13.5.5329.00 port: 49,572 cefVersion: 0 name: A management user has successfully authenti
cated deviceAddress: 172.16.19.250 sourceAddress: 172.16.10.156 address: 172.16.19.240 deviceEventClassId: 352
```

Limitations/considerations

The following limitations need to be considered:

- Each cluster node needs to have a public management address to be able to communicate with the Logstash server.
- Syslog events can be lost when the NAS server system reboots, due to network unavailability. They are also sent via UDP, so onward transmission by the network is not guaranteed.
- When using a cluster, the events are not necessarily sent from the same address, even with the `syslog_and_snmptrap_come_from_admin_vnode` variable set. The admin EVS sends most events but can select another serving EVS address if there is no route from the admin EVS to the Logstash server.
- Auditing events come from both physical node addresses, rather than specific admin address, as they are specific to the individual node.

Note: The use of multiple addresses makes filtering for events from a complete system more difficult, as they may have come from any of the configured IP addresses.

Hitachi Vantara

Corporate Headquarters 2535 Augustine Drive
Santa Clara, CA 95054 USA www.HitachiVantara.com community.HitachiVantara.com

Regional Contact Information

Americas: +1 866 374 5822 or info@hitachivantara.com

Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com

Asia Pacific: +852 3189 7900 or info.marketing.apac@hitachivantara.com

